

LISE pre equilibrium research development

Complete physics and software record through the temporary research freeze

17 October 2026 | Source snapshot 18.8.91 | Pre equilibrium milestone 18.8.64

Purpose and current conclusion

This report consolidates the two earlier LISE pre-equilibrium Word documents and the subsequent development into one technical reference. It explains why the project began with thermalization, how it changed to a forward calculation, what the kernel means, how probabilities and conservation are implemented, and which numerical and interface decisions were retained.

The implemented product is a probability-based research framework for abrasion prefragment → fast neutron and proton emission → existing LISE statistical de-excitation. Identity and algebraic supplied kernels validate the software. A predictive physical fast-emission model has not been adopted. The research path has not replaced ordinary LISE abrasion–ablation production calculations.

Current operating decision. Development is temporarily frozen. The public controls in the Evaporation Calculator and Excitation Energy dialog remain visible but disabled. Their master state is forced to Off when these dialogs initialize or refresh. Research modules, configuration metadata and direct tests remain available in the source for a later restart.

Reading this record

Sections 1–6 define the physics and probability architecture. Sections 7–12 document implementation, numerical decisions and exchange interfaces. Sections 13–17 record performance, validation, user workflow and restart conditions. The appendices provide a source map, version history and evidence register.

Evidence boundary. Timings and regression outcomes below are historical measurements recorded during development, principally through 10 October 2026. Preparing this document did not rebuild LISE or rerun its physics tests. Application version 18.8.91 is the source snapshot on 17 October; later unrelated application changes are outside this report.

Contents

- 1 Development scope and evolution
 - 2 Thermalization and the legacy calculation
 - 3 Main forward architecture
 - 4 Abrasion domains and raw excitation
 - 5 Kernel definition and conservation
 - 6 Probability quantities and final state support
 - 7 Worked probability example
 - 8 Kernel validation and application scheme
 - 9 Statistical decay integration and channel meaning
 - 10 Numerical development and retained backend
 - 11 Runtime safeguards and numerical diagnostics
 - 12 Supplied kernel and provider exchange
 - 13 File schema persistence and result isolation
 - 14 Synthetic validation fixtures and observed output
 - 15 Measured scaling for calcium lead and uranium
 - 16 Validation evidence and acceptance criteria
 - 17 User interface freeze and restart roadmap
- Appendix A Source map for future development
- Appendix B Development chronology and detailed version ledger
- Appendix C Evidence register and references

The Word Navigation Pane lists the report sections and version entries. Figures 1 to 3 show the main forward path, kernel procedure and external-provider exchange. Tables distinguish probability definitions, benchmark results, implementation responsibilities and historical milestones.

1 Development scope and evolution

The initial question was whether LISE could add an optional neutron/proton pre-equilibrium stage while retaining its existing abrasion excitation and statistical-decay infrastructure. ABRABLA07 was examined only for comparison; no ABRABLA emission physics, fitted parameters, tables or code were adopted.

From backward inference to forward propagation

The first research design selected a final isotope and used the existing full prefragment search to identify possible abrasion ancestors. This remains useful for interpreting parents conditioned on a rare residue, but it made the overall scheme look as if the final fragment had to be known before the reaction could be calculated. It also raised concerns about heavy-system runtime.

The adopted main architecture instead starts with one abrasion prefragment, or a declared forward ensemble of abrasion prefragments. Each parent passes through the kernel and existing statistical decay to produce a final distribution. A selected isotope is used afterward to condition that distribution and calculate its parent and fast-emission posterior. The backward search is retained only as an optional diagnostic.

Requirements retained throughout

- Default Off preserves the legacy calculation and its original outputs. Research results contain dimensionless probabilities and support, with energy diagnostics in MeV.
- Fast emission and subsequent equilibrium statistical decay remain separate stages. `SCelement::Widths` supplies downstream statistical physics, never a fast-emission kernel by assumption.
- Every supplied branch must conserve mass number, charge and the declared energy ledger. Missing or invalid input blocks the calculation with an explicit error.
- Rare proton-selective parents must not disappear through an N/Z -direction restriction or a hidden top-parent cutoff.
- The no-pickup research scope and $L = 0$ remain explicit. The selected forward abrasion depth defines a declared domain rather than a hidden physics approximation.

Three different meanings of frozen

The version-1 interchange contract is frozen as a stable schema. The project itself is paused by the user. The temporary UI gate disables research access in the two public dialogs. These are separate decisions: a stable file schema does not imply that a physical model is finalized, and disabling the interface does not remove the validated research implementation.

2 Thermalization and the legacy calculation

The inspected implementation follows Gaimard–Schmidt Eq. 3.4 through the dimensionless factor $F(q)$. The source explicitly identifies this factor as the unthermalized fraction. The following equations describe the implemented LISE treatment; they do not define an emission law.

$$F(q) = \int_0^1 \exp[-q(1 - u^2)] du \quad q = (E_0^*/A_0) t_{abr} / (2\alpha)$$

`CmThermalizationFactor` in `o_Fireball.cpp` constructs the abrasion time from a geometrical path and projectile velocity. The present path length uses $\max(C_p, C_t) + 0.474 \min(C_p, C_t)$; α is obtained from the configured thermalization time coefficient with the existing bounds. These are existing LISE settings and geometry, not new kernel parameters.

`Thermalization::fraction` uses a positive convergent series at $q < 40$ and an endpoint asymptotic expansion at larger q . Nonpositive q returns one and nonfinite positive input tends to zero. The thermalization development also addressed double application, stale excitation distributions and conservative remapping when the transformed energy axis folds.

What LISE passes onward

Legacy excitation transformation $E_0^* \rightarrow E_0^* F(q)$

`Fireball::ApplyThermalization` documents this as phenomenological suppression of the diabatic hole excitation supplied to evaporation. Eq. 3.4 describes equilibration; interpreting the removed excitation as energy carried by fast particles would require another physical assumption. The current legacy suppression does not assign $E_0^*[1 - F(q)]$ to explicit neutron/proton separation, kinetic or recoil energies.

Order of existing processing

Abrasion excitation → optional thermalization suppression → conditional temperature limiter → statistical decay


The limiter maps temperatures above the existing threshold through $T_{new} = T_{lim} + 0.03(T - T_{lim})$, then converts temperature back to excitation. The legacy implementation includes distribution remapping and broadening. The limiter is not a microscopic fast-particle energy ledger.

Research decision

Research preparation takes raw abrasion excitation from an isolated Fireball geometry with thermalization disabled. The research path bypasses the legacy suppression and the limiting-temperature remapping. $F(q)$ is reported only as context, usually at the raw mean. Neither $F(q)$ nor $1 - F(q)$ is interpreted as an emission probability. Identity therefore tests direct decay of identical raw excitation; it need not reproduce a legacy run with suppression or limiting temperature enabled.

The original-paper interpretation remains a prerequisite for any future physical connection between equilibration and an emission kernel. The implementation comments settle what LISE currently computes; they do not establish a new predictive dynamical model.

3 Main forward architecture



Physical sequence runs forward. Final isotope selection is an analysis after propagation.

Figure 1 Adopted architecture of the retained research implementation. All statistical channels are downstream of the same pre-equilibrium stage.

Single prefragment

The selected-parent calculation uses one (A, Z) with prior one, prepares its normalized raw excitation distribution and produces all retained final residues plus fission, LISE statistical breakup and numerical accounting. It is the smallest useful validation and runtime case.

Forward ensemble

The ensemble enumerates every physically allowed neutron/proton composition for removal counts 1 through the declared maximum. Geometry and composition weights become normalized dimensionless abrasion priors. For three removed nucleons, the Ca/Pb/U benchmark domains contain nine parents. No final-isotope N/Z trajectory selects those parents.

Conditioning

The chosen final isotope does not generate the forward ensemble. It selects one support value from the computed distribution, then normalizes the corresponding joint contributions to obtain a posterior. A different final isotope changes the conditioning; comparisons require a common parent domain, prior definition and settings.

4 Abrasion domains and raw excitation

Forward parent domain

ForwardEnsembleLISE::makeDomain iterates removed nucleons $d = 1 \dots d_{\max}$ and removed protons $y_{abr} = 0 \dots d$, with removed neutrons $x_{abr} = d - y_{abr}$. It checks projectile proton/neutron availability and valid residual charge. The unchanged projectile is excluded. Here abrasion multiplicities label removal before the kernel; they must not be confused with fast multiplicities x and y within the kernel.

$$w_i = \text{geometrical weight} \times \text{composition weight} \quad \pi_i = w_i / \sum_j w_j$$

GetGeomCrossSection and AAfactorial provide existing LISE geometrical and composition weights. Their internal origin does not change the meaning of the result: after normalization, π is dimensionless within the declared domain. Research reports do not save or display an absolute cross section.

Backward diagnostic domain

The optional backward calculation uses the full candidate scan underlying TEvaporationDlg::CmPlotInitPrefragment. Candidate proton and neutron counts span the selected final isotope through the projectile, excluding the unchanged projectile. There is no N/Z -distance or top-parent cutoff. This preserves rare proton-selective ancestors such as ^{55}Cl and ^{55}S for ^{46}Si whenever projectile composition and nuclear data permit them.

The forward depth-limited domain and a final-conditioned full backward domain are generally different ensembles. Their L_f values must not be compared as if they had identical normalization. The report must state which domain and priors were used.

Raw excitation distribution

$$\int p_i(E_0^*) dE_0^* = 1 \quad a_{ik} \geq 0 \quad \sum_k a_{ik} = 1$$

Raw excitation is obtained from the configured LISE abrasion excitation model. It is not replaced with one common temperature for every isotope. The supplied interface uses exact energy nodes and normalized nodal probability masses computed by trapezoidal quadrature of a nonnegative piecewise-linear density. The main Identity route uses the existing normalized LISE excitation-grid representation.

Normalization of the initial ensemble is intentional. Discarding failed kernel branches and renormalizing the surviving branches is prohibited. Zero prior does not authorize redistribution; invalid or unavailable data on required positive-prior paths produce an explicit failure. A failed or cancelled scan never becomes a complete posterior.

5 Kernel definition and conservation

The kernel is the conditional distribution of fast neutron/proton outcomes and residual excitation for one abrasion parent at one raw excitation. It is independent of which final isotope will later be selected for analysis.

“Unconditional” kernel reports mean unconditional on that final isotope; the kernel remains conditional on parent and input excitation.

$$\Sigma_{xy} \int K_i(x,y,dE_r^* | E_0^*) = 1$$

A discrete branch carries an identifier, neutron multiplicity x , proton multiplicity y , probability mass k , residual excitation E_r , total fast-particle kinetic energy K_{fast} and recoil energy K_{recoil} . Multiple branches with identical multiplicities may carry different residual excitations. This correlation must survive aggregation.

$$A_r = A_i - x - y \quad Z_r = Z_i - y$$

$$\Delta M = M_{nucl}(A_r, Z_r) + x m_n + y m_p - M_{nucl}(A_i, Z_i)$$

$$E_0^* = E_r^* + \Delta M c^2 + K_{fast} + K_{recoil}$$

M denotes a ground-state bare nuclear mass under the consistent LISE convention, and m_n and m_p are free-neutron and free-proton masses. The implementation supplies rest energies in MeV using `Celement::Mnucl() × amu_MeV` with the configured LISE electron-binding correction. Atomic masses cannot be inserted as bare nuclear masses.

$\Delta M c^2$ already represents the ground-state mass/separation contribution. A Coulomb barrier must not be added a second time as a separate energy loss. The kernel provider may use a barrier in an emission-rate model, but its final residual, kinetic and recoil values must still satisfy this one ledger.

Identity operator

$$K_{00} = \delta(E_r^* - E_0^*) \quad A_r = A_i \quad Z_r = Z_i \quad K_{fast} = K_{recoil} = 0$$

Identity is exact as a zero-emission operator on the chosen input representation. It does not claim that the input quadrature or downstream statistical calculation is continuum-exact. The difference between the raw Identity route and suppressed legacy excitation is deliberate and must not be interpreted as a physical pre-equilibrium signal.

Scope of conservation checking

The scalar energy ledger checks total energy, not momentum vectors or angular correlations. Recoil provenance must explain how the provider generated recoil. A new momentum-vector, cluster-emission or spin/parity representation requires a new interchange version. No later temperature transformation may silently remove energy from a research residual.

6 Probability quantities and final state support

Quantity	Meaning and normalization
π_i	Abrasion prior normalized across the declared parent domain.
a_{ik}	Raw excitation-node probability for parent i; sums to one.
$k_{ik\beta}$	Branch probability given parent i and input node k; sums to one at that node.
$U_i(x,y)$	Input-energy-averaged fast multiplicity probability before selecting final f.
$Df(A_r, Z_r, E_r^*)$	Existing statistical probability of reaching f from a residual state.
$W_{ixy}^{(f)}$	Combined contribution to f before conditioning; dimensionless.
L_f	Sum of W for f within the normalized declared ensemble.
$P(i,x,y f)$	Posterior given f, defined only for positive L_f .
Omitted probability	Explicit computational mass loss; never redistributed.
Numerical excess	Positive balance excess reported separately; not hidden in omitted mass.

$$U_i(x,y) = \sum_k a_{ik} \sum_{\beta \in (x,y)} k_{ik\beta}$$

$$W_{ixy}^{(f)} = \pi_i \sum_k a_{ik} \sum_{\beta \in (x,y)} k_{ik\beta} Df(A_{r,\beta}, Z_{r,\beta}, E_{r,\beta}^*)$$

$$L_f = \sum_{ixy} W_{ixy}^{(f)} \quad P(i,x,y | f) = W_{ixy}^{(f)} / L_f$$

Parent and multiplicity marginals are obtained by summing the joint posterior: $P(i | f) = \sum_{xy} P(i,x,y | f)$ and $P(x,y | f) = \sum_i P(i,x,y | f)$. The unconditioned ensemble multiplicity is $U(x,y) = \sum_i \pi_i U_i(x,y)$. It is generally different from $P(x,y | f)$.

For a grouped row with $\pi_i U_i > 0$, the reported energy-averaged statistical D is $W/(\pi_i U_i)$. It is a conditional average over that row's input/kernel residual states. Multiplying π_i , U_i and that D reconstructs W. A generic unweighted average of decay probabilities does not.

L_f measures model support within a normalized, explicitly declared domain. It is neither an absolute all-collision production probability nor a cross section. Do not renormalize L_f across selected isotopes. If $L_f = 0$, report “no supported path” and an undefined posterior.

7 Worked probability example

The following algebraic example illustrates the reporting contract. It is not a physical calculation. Two abrasion parents have priors 0.6 and 0.4. After input-energy averaging, each has Identity and 1n branches. The statistical column is the probability of reaching one selected final isotope from that branch's residual ensemble.

Parent	π_i	x y	U_i	Df	W	Posterior
1	.60	0 0	.75	.10	.045	.225
1	.60	1 0	.25	.50	.075	.375
2	.40	0 0	.50	.20	.040	.200
2	.40	1 0	.50	.20	.040	.200

$$L_f = 0.045 + 0.075 + 0.040 + 0.040 = 0.200$$

$$P(1 | f) = 0.600 \quad P(2 | f) = 0.400$$

$$U(1,0) = 0.6 \times 0.25 + 0.4 \times 0.50 = 0.350$$

$$P(1,0 | f) = 0.375 + 0.200 = 0.575$$

The final isotope favors the fast-neutron branch: its conditional multiplicity probability is 0.575 although the unconditioned multiplicity probability is 0.350. This does not change the input kernel; it reflects selection by the subsequent decay probability.

Why support must accompany a normalized posterior

Halving every D in this example halves every W and gives $L_f = 0.100$. All posterior entries and marginals remain identical. A normalized posterior alone would hide this factor-of-two difference in final-state support.


Zero support and numerical loss

If all four statistical probabilities are zero, L_f is zero and the posterior is undefined. One must not substitute a uniform posterior or remove unsupported parents and renormalize their priors. A competing physical residue, fission or statistical breakup is an outcome, not numerical loss.

$$\sum L_f + P_{\text{fission}} + P_{\text{breakup}} + P_{\text{omitted}} - P_{\text{excess}} \approx 1$$

This is a balance check over a complete retained residue distribution and the separately recorded terminal channels. P_{excess} is a diagnostic that explains a small overbalance; it does not authorize cancelling an excessive probability against omitted mass. Bounds and tolerance checks remain required. Probability closure does not bound quadrature or physics error for an individual rare isotope.

8 Kernel validation and application scheme



Invalid tables are rejected. No interpolation, extrapolation, repair or survivor renormalization.

Figure 2 Complete validation precedes application. The failure block is an alternative exit, not a physical branch.

Strict acceptance rules

Required records must cover every positive-prior parent and every exact required input-energy node. The reader rejects unknown fields, duplicate JSON keys including escaped duplicate spellings, invalid nuclei/masses, nonfinite or negative values, incorrect branch sums and energy-ledger residuals outside tolerance. Multiple branches may share multiplicity while retaining distinct energies.

Node normalization uses absolute tolerance 10^{-8} . The energy ledger uses $10^{-8}(1 + E_0)$ MeV. Accepted values are not repaired or renormalized. Matching-grid checks are separate from the energy-ledger tolerance; satisfying conservation does not permit a different input grid.

The supplied file is pinned by SHA-256 over its exact bytes. Selection validates file structure; calculation validates the full reaction and mass context. Changed file bytes require explicit reselection. Changed physics settings require a matching table. These errors must never silently substitute Identity or the legacy calculation.

9 Statistical decay integration and channel meaning

The new stage ends at a residual nucleus and residual excitation. Existing LISE statistical de-excitation then evaluates the enabled competing channels. Ablation through statistical particle emission, fission and LISE statistical breakup are downstream alternatives. Statistical breakup is the pre-existing LISE channel; it is not a synonym for the new pre-equilibrium stage.

Continuous Identity input

The research adapter supplies a unit-integral excitation density to ManagementAA, with the existing AbraAbla statistical machinery downstream. It uses probability population one and avoids the legacy result object, cross-section plot distributions and cross-section output files. The Identity path can therefore be validated against direct statistical decay of the same raw density.

Discrete supplied residual input

A branch at $E_r = 0$ is a point mass, not a narrow continuous distribution. A nonnegative continuous density on $E \geq 0$ has a positive mean unless represented as a true point. Narrowing a distribution would alter the supplied residual energy and introduce grid dependence. The discrete adapter evaluates the first residual excitation exactly.

In 2T mode, the subsequent chain follows the existing mean-energy daughter law recursively. In qualitative mode, it uses the existing $(U - E)p(E)$ law and passes the daughter density to ManagementAA. Exact first-state input does not make every later statistical approximation exact; grid convergence of downstream densities remains necessary.

Validated special treatment

A completely validated all-zero-emission unit table can be evaluated as Identity on the raw continuous density. Generated synthetic validation tables also have a bounded evaluator that preserves the Identity part of the normalized input density while treating emitted residuals at their declared energy. This avoids needlessly expanding a high-energy Identity distribution into expensive point chains.

General supplied kernels use their actual residual-state branches and the general point-decay adapter. The synthetic dispatch is a software-validation route, not permission to treat an arbitrary physical table as one of the known fixtures. Complete validation remains required before any optimization.

Explicit limits

Research calculations reject excitation beyond the width solver's $T = 20$ MeV numerical range rather than silently clamping energy. A zero-energy unbound nucleus needing prompt decay is rejected pending an exact treatment. The discrete-state ceiling reports failure rather than a completed truncated distribution. Missing masses, cancellation and invalid statistical probabilities likewise prevent publication of a completed result.

10 Numerical development and retained backend

The large numerical investigation was needed because conservation and normalization did not establish the accuracy of rare final-isotope support. The original architecture-and-status document at 18.8.21 captured an intermediate stage of this investigation; several items then listed as next steps were subsequently completed.

Development phase	Finding and consequence
Point chains and memoization	High-energy exact trees could be costly; exact-state memoization had zero hits in the measured individual trees.
Conservative cascade	Probability and mean excitation were merged within nucleus/energy bins. Thirty Re/Pb/U cases completed, but binning and input quadrature needed independent checks.
Pb phase and input studies	Transport accuracy and input quadrature were isolated. Adaptive seed dependence exposed a false impression of convergence.
Support and value jumps	Support-only discovery missed transfers between already-present isotopes. Value-jump detection and corrected acceptance gates were added.
U steep interval	A steep continuous residue/fission transfer was localized near 28.338615 MeV and resolved with independent dense and adaptive comparisons.
Reusable adaptive integrator	Transition discovery and adaptive probability integration became reusable research components; unresolved input mass stayed explicit.
Practical Identity decision	Si-30 discovery remained too costly. Main user-facing Identity returned to the bounded normalized LISE grid; adaptive integration remained diagnostic.

Measured numerical evidence

The 18.8.21 record reported an adaptive starting-grid spread around 5×10^{-7} despite an indicator of 10^{-7} . Later transition-split Pb studies reduced a finest step difference to 2.21×10^{-10} isotope L1 in the tested cases. Correcting the unmatched Au bracket gave matching 96-bracket maps at two tested floors. These are measured consistency results, not certified global accuracy bounds.

For U-236, a localized dense reference reached 1.06×10^{-13} isotope L1 in the studied interval; vector adaptive Simpson agreed within 3.10×10^{-13} . Full-domain tightening reduced cross-grid isotope L1 from 1.44×10^{-8} to 1.19×10^{-9} . The measured success on those inputs did not eliminate the Si-30 runtime problem or prove completeness of finite sampling.

11 Runtime safeguards and numerical diagnostics

Current practical route

From 18.8.35 onward, the selected-prefragment Identity action uses the normalized LISE excitation grid and existing ManagementAA decay. The high-resolution adaptive point integrator remains a separately identified numerical diagnostic. The conservative cascade and alternative traversals remain preserved prototypes, not a hidden replacement for production decay.

What is counted

- Input nodes and branch records describe kernel input size. Positive records and emitted records are distinguished before decay.
- Exact within-parent merging reduces repeated identical (x,y,A,Z,Er) states. Decay caching can share identical (A,Z,Er) maps across multiplicities and parents.
- Continuous Identity distributions remain parent-local and do not collide with point-state cache entries, including $Er = 0$.
- Preparation, statistical and total time, actual decay evaluations, cache hits and calculated nuclei are reported. Runtime per executed decay is an observational metric.

Pre run workload forecast

The inspector counts merged emitted states, globally unique emitted residual states and potential cross-parent reuse without statistical decay. It explicitly excludes parent-local continuous Identity work. This forecast helps estimate costs before increasing the abrasion domain, but it is not a guaranteed wall-clock prediction.

Cancellation and publication


Ensemble progress maps preparation and each parent to one monotonic progress sequence. Inner decay loops pump cancellation checks. Cancellation, a resource ceiling or an adapter exception produces “No completed result”; export is not enabled for a partial scan. A late configuration change or changed pinned file must invalidate a result before publication.

Interpretation of error measures

Omitted probability measures explicitly identified computational mass deficits. Numerical excess is recorded independently. Adaptive indicators estimate integration behavior in the sampled regions. Grid refinement and independent-reference isotope L1 comparisons measure differences between calculations. None of these alone is a rigorous total error bound on rare-isotope L_f , and none validates the physical emission model.

Isotope L1 is the sum over isotopes of absolute support differences, without posterior renormalization. It can be small while a very rare isotope has a large relative change. A complete validation report should retain both aggregate differences and the supports of scientifically relevant rare residues.

12 Supplied kernel and provider exchange



Request is input data, not an emission kernel. A valid identity response proves exchange plumbing only.

Figure 3 Separation of LISE data preparation, external physics, validation, explicit selection and downstream statistical calculation. UI access is currently frozen.

LISE-pre-equilibrium-request-1 is a companion provider package. It includes normalized parents, exact excitation-node masses, reaction and mass signatures, free nucleon rest energies, and all LISE nuclear masses reachable through neutron/proton removal. Unavailable nuclei are listed explicitly. It contains no emission probabilities, residual-energy law or multiplicity cutoff.

The response is a separate LISE-pre-equilibrium-1 kernel. KernelProviderExchange checks request completeness and reconstructs the exact context before validating the response. The console LISEKernelValidator and the UI request/response action share this contract. Validation does not select the response, change a completed calculation or run statistical decay.

LISEKernelValidator provider-request.json kernel-response.json

Console exit codes are 0 for a valid exchange, 1 for invalid data or processing failure, and 2 for incorrect arguments. Files are subject to the tool's declared size checks. A valid exported Identity template is not evidence for a nonzero physical fast-emission model.

13 File schema persistence and result isolation

Record	Required data
Kernel root	format, name, version, provenance, recoil_provenance, units, frame, reaction_signature, mass_signature, parents
Parent	A, Z, nodes
Input node	E0, branches in strictly increasing matching input-energy order
Branch	id, x, y, probability, Er, Kfast, Krecoil
Units and frame	MeV and abrasion-parent-rest
Configuration	PreEquilibriumEnabled, PreEqKernelType, PreEqKernelVersion, PreEqKernelProvenance, PreEqKernelFile, PreEqKernelDigest

Configuration lifecycle

Complete configuration loads reset research settings before parsing. Older files without fields therefore start Off even after a preceding On session. Relative table paths resolve against the configuration location. Inactive kernel settings remain stored when Off without loading or validating the file.

Metadata uses the existing serializer through PEQ64 UTF-8/Base64 encoding so Unicode, multiline text and comment characters survive round trips; older plain-text fields remain readable. Encoded metadata/path fields have an explicit 12,000-byte bound. The exact byte checksum binds a selected file, and a content change requires explicit reselection.

Separate research results

Research results and caches do not populate legacy cross-section caches or save sections. Changing mode, kernel, relevant reaction/physics settings, parent domain or conditioned isotope invalidates the displayed research result and disables export until recalculation. Result export writes a text report containing probabilities, domain, metadata and diagnostic quantities.

Frozen interface behavior

The temporary release gate `researchUserInterfaceAvailable` is false in `o_PreEquilibriumSettings.h`. The gate changes the live enabled value to zero when either public dialog initializes or refreshes. It preserves inactive kernel metadata. This is a UI freeze, not a change to the underlying supplied-file schema or the mathematical kernel APIs.

The gate is not applied by the configuration parser itself. A loaded file can retain its parsed On request until one of the guarded UI lifecycle methods enforces Off. Direct research APIs and test harnesses remain callable with explicit research settings. Therefore the freeze should be described precisely as disabled public dialog access, not as removal of every research entry point from the executable.

14 Synthetic validation fixtures and observed output

Three generated fixture families exercise the supplied route without claiming an emission parameterization. Each is labeled synthetic and carries version/provenance information. Threshold-inaccessible branch probability remains in Identity when the fixture is constructed; a general supplied table is never repaired after loading.

Fixture	Purpose
One neutron	Algebraic $P(1n) = 0.200$ above threshold; residual $E_r = 0$ and two-body kinetic/recoil partition.
Correlated n and p	Algebraic $P(1n) = 0.100$, $P(1p) = 0.100$ and $P(1n1p) = 0.050$ when allowed. Tests multiplicity correlation and posterior marginals.
Finite residual excitation	Same n/p families with declared emitted residual excitation 40.000 MeV when energetically allowed. Tests nonzero- E_r decay and cache sharing.

Reported 42S example

In the recorded user run for selected $A = 42$, $Z = 16$, Identity raw mean excitation was 143.90 MeV and $F(q)$ at the raw mean was 0.702. The following probability values are retained historical output for that particular configuration, not a reference result for every 48Ca reaction.

Outcome	Identity	Synthetic one neutron
Total residue probability	0.99862	0.99890
Fission	1.2415e-03	9.9319e-04
LISE statistical breakup	1.3907e-04	1.1126e-04
Omitted numerical probability	1.1047e-08	8.8166e-09
Numerical excess	2.3083e-15	1.1621e-15

The synthetic 1n branch accounted for approximately 0.2 of the probability while Identity retained approximately 0.8. The example is consistent with a zero-residual-excitation emitted branch contributing directly to the residual isotope. The figures demonstrate software propagation and accounting; they are not evidence that a physical pre-equilibrium probability is twenty percent.

Readable output

The UI formats raw mean excitation with two decimals, $F(q)$ with three, total residue probability with five, and small channel probabilities, omissions and excesses with four decimal scientific notation. Isotope probability rows use the same scientific format. Runtime and provenance parameters use compact formatting; stored values retain full precision.

15 Measured scaling for calcium lead and uranium

The 10 October 2026 benchmark used carbon, 1 GeV/u projectiles and 64 excitation intervals. Identity and the synthetic finite-Er kernel were run at abrasion depths one through three. The table reproduces recorded single-run engineering timings; they are not performance guarantees for a new physical kernel.

Beam	Max removal	Parents	Identity	Synthetic	Decays	Reuse
48Ca	1	2	68 ms	76 ms	7	12.50%
48Ca	2	5	350 ms	398 ms	15	25.00%
48Ca	3	9	1.17 s	1.37 s	25	30.56%
208Pb	1	2	104 ms	161 ms	7	12.50%
208Pb	2	5	811 ms	830 ms	15	25.00%
208Pb	3	9	2.53 s	2.90 s	25	30.56%
238U	1	2	212 ms	236 ms	7	12.50%
238U	2	5	1.20 s	1.28 s	15	25.00%
238U	3	9	3.13 s	3.16 s	25	30.56%

All 18 cases satisfied the declared probability balance tolerance and constructed a defined posterior for the residue of strongest support. Decay counts and reuse refer to the synthetic calculation; parent-local Identity distributions are not point-state cache entries.

Provider package inventory

Projectile	Reachable nuclei	Available masses	Package size
48Ca	607	607	146 KiB
208Pb	10,539	10,539	1.17 MiB
238U	13,669	13,669	1.50 MiB

These inventories cover the union of nuclei reachable by neutron/proton removal from the declared nine-parent domain, not merely the nine abrasion parents. All reachable nuclei in these measured packages had LISE masses. Export was cancellable and used atomic file replacement.

Runtime implications

The tests show practical shallow-domain operation for the measured heavy beams. Cost can increase with abrasion depth, input-grid density, distinct residual energies and statistical settings. Continuous physical residual-energy laws represented by many discrete nodes can have far fewer cache hits than the finite-Er fixture. No parent cutoff or hidden energy merging is justified by these timings.

16 Validation evidence and acceptance criteria

Validation was layered to distinguish probability algebra, file handling, statistical integration, interface behavior and physical validity. The historical notes repeatedly record all twelve exact legacy reference comparisons passing during the relevant production runs. That coverage is a tested set, not proof for every reaction and configuration.

Layer	Checks and scope
Pure probability	Analytic W, Lf and posterior; marginals; identical posterior with different support; zero support; probability bounds.
Kernel contract	Schema, duplicate keys, metadata pin, complete grid/domain, normalization, mass availability and energy conservation; explicit rejection without repair.
Production adapter	Actual LISE masses and statistical decay, Identity equivalence, synthetic n/p paths, exact first residual energy, heavy-nucleus and grid checks.
Legacy isolation	Off bypasses research processing; historical abrasion/decay cases compare exactly; separate result/cache ownership.
Provider exchange	Request inventory, exact context, response validity, console/UI round trip, cancellation and no automatic kernel selection.
Qt interface	Defaults, invalid requested kernels, settings persistence, separate dialog, legacy controls and temporary freeze.
Runtime and numerics	Ca/Pb/U scaling, cross-parent cache use, monotonic progress, cancellation, input/grid and reference comparisons.

What remains unproved

The scalar ledger does not verify a full momentum distribution. Finite grids do not guarantee rare-isotope accuracy. Normalized branches do not establish a physical emission rate. A supplied table with placeholder provenance may be structurally valid while representing Identity only. Synthetic fixtures validate code, not an experimental prediction.

Checks to repeat on restart

After reopening the gate, rebuild affected production objects and run the supplied-kernel, provider-package/exchange, focused Qt and production integration suites appropriate to the change. Recheck legacy regression, complete result invalidation, explicit error handling and no export after cancellation. For a proposed physical kernel, add independent physics validation and measured convergence on the relevant rare isotopes.

The latest freeze change at 18.8.64 compiled both affected dialog objects and passed the focused temporary-freeze and settings tests. The complete application was not linked by that change alone. This report records those outcomes without claiming a new test execution on 17 October.

17 User interface freeze and restart roadmap

Public state at the pause

The compact master/kernel/settings/launch strip was moved near the bottom of the Evaporation Calculator at 18.8.63. The separate modeless research window introduced at 18.8.62 has Calculations, Kernel exchange and Software validation tabs. At 18.8.64, the public controls were disabled at the user's request while remaining visible. The Excitation Energy checkbox and Settings button are also disabled.

The research implementation is preserved for future work. The gate is centralized in `o_PreEquilibriumSettings.h`. Reopening it is a source change requiring validation; no public On option is presently intended for normal operation.

Restart in a controlled order

- Confirm the temporary pause has ended and review changes made to LISE statistical or configuration code during the pause.
- Reopen public access through the centralized gate, rebuild and repeat focused UI plus legacy isolation checks.
- Repeat a selected-prefragment Identity calculation, then a known synthetic supplied table and a small forward ensemble.
- Select a defensible physical model or external provider and specify the non-equilibrium initial state, transition/emission rates, stopping criterion and applicability.
- Generate a reaction-matched response, validate conservation and signatures, then compare single-parent U and final supports with Identity.
- Extend to declared Ca/Pb/U domains and scientifically relevant residues, reporting support, posteriors, omission, convergence and runtime.
- Consider ordinary AA integration only after a separate physics and product decision. Maintain exact default-Off behavior and an explicit experimental label.

Outstanding physics inputs

The current abrasion state contains averaged (A, Z, E^*) information. It does not uniquely determine an initial exciton particle-hole configuration or an eventwise spatial nucleon arrangement. Choosing an exciton or clustering label without supplying those required inputs would introduce an unvalidated assumption. The provider interface is ready to receive a justified external calculation, but it cannot supply missing physics itself.

A candidate provider must document the initial proton/neutron state, rates and parameter provenance, the fast-to-statistical stopping rule, correlated multiplicities, residual excitation, kinetic-energy distribution, recoil prescription and validation data. No ABRABLA07 input is permitted. Temperature should follow each actual residual state and LISE settings rather than an imposed common temperature for all isotopes.

A Source map for future development

Source or module	Responsibility
o_Evap/o_Thermalization.h	Stable $F(q)$ evaluation and conservative remapping helper.
o_Evap/o_Fireball.cpp	Raw abrasion excitation, thermalization and legacy temperature transformation.
o_Evap/o_PreEquilibrium.h and .cpp	Nucleus/outcome definitions, candidates, normalized priors, contribution and posterior algebra.
o_Evap/o_PreEquilibriumSettings.h	Persistent settings defaults and temporary public-interface availability gate.
o_Evap/o_PreEquilibriumLISE.cpp	Existing-statistical adapter and optional full backward Identity analysis.
ForwardPrototype / ForwardPointDecay	Raw normalized single-parent Identity preparation and statistical response.
ForwardLISE / ForwardSupplied	Single-parent dispatch, supplied residual propagation and reporting.
ForwardEngine / ForwardEnsemble	Reusable state/decay maps, exact merging/cache, ensemble combination and conditioning.
ForwardEnsembleLISE	Forward domain preparation, priors, ensemble dispatch, timing and report integration.
SuppliedKernel / SuppliedLISE	Strict JSON contract, pin, context/mass validation, LISE matching grids and signatures.
DiscreteDecay	Exact declared initial residual excitation and existing LISE downstream laws.
ForwardSyntheticValidation	Bounded evaluator for generated synthetic validation fixture families.
KernelProviderPackage / Exchange	Complete input/mass inventory export and strict request/response validation.
ResearchProbabilityIntegrator	Reusable probability-transition discovery and adaptive integration diagnostic.
ConservativeCascade / ConservativeLISE	Preserved conservative transport prototype; not default application decay.
ResearchDialogUi / ResearchActions / ProviderExchangeUi	Grouped dedicated research window and kernel/provider actions.
d_Evaporation_research.cpp	Compact master strip, research window hooks, freeze synchronization and results.
d_Apf_excitation.cpp / d_PreEquilibriumSettingsDialog	Excitation dialog controls and separate kernel settings transaction.
c_ConfigPP1.h / c_ConfigPP_read.cpp	Serialized fields and reset before complete configuration load.

Module names without a full path above are under o_Evap/pre_equilibrium unless explicitly qualified. ManagementAA / AbraAbla and SCelement remain the existing downstream statistical machinery.

B Development chronology

The chronology below separates the pre-equilibrium milestones from unrelated application changes. Dates and version associations follow `L_Init/lise_version.h`. The detailed version ledger continues afterward so work can be resumed without reconstructing the conversation.

Period and version	Development result
15 Sep 18.8.12–13	Thermalization numerical/call-path corrections and availability; documented the additional LISE suppression assumption.
16–18 Sep 18.8.14–18	Stage 1, configuration isolation, settings UI, strict supplied kernel and forward architecture.
25 Sep–1 Oct 18.8.19–21	Probability accounting, transport comparisons, conservative cascade and independent input quadrature.
2–5 Oct 18.8.22–31	Pb/U interval diagnosis, support/value transitions, corrected discovery gates and adaptive references.
5–6 Oct 18.8.32–35	Reusable research integration; cancellation and Si-30 performance; retained bounded Identity user path.
7 Oct 18.8.36–51	Readable reports, UI sizing/reflow, synthetic 1n and n/p fixtures, finite residual excitation, inspector and guide.
8–10 Oct 18.8.52–56	Ensemble metrics, shared emitted-state cache, cancellation/progress, workload forecast and scaling metrics.
10 Oct 18.8.57–61	Ca/Pb/U benchmark and contract decision; provider packages, exchange validator, console and UI validation.
10 Oct 18.8.62–64	Dedicated grouped dialog; compact strip moved to bottom; user-requested visible but disabled freeze.
17 Oct source 18.8.91	Documentation consolidation; research gate still false. Later unrelated LISE changes are outside scope.

How the two earlier documents are incorporated

`LISE_pre_equilibrium_schemes.docx` supplies the original kernel conservation and probability definitions. Its final-selected global scheme is replaced here by the adopted forward scheme.

`LISE_PreEquilibrium_Architecture_and_Status_18_8_21.docx` supplies the forward architecture and numerical status as of 1 October; its planned investigations are updated here with the later results, retained bounded backend and interface pause.

B1 Detailed version ledger

18.8.12 09/15/26

Abrasion thermalization correction Prevented double application in the Gaimard exponential fallback and stale excitation distributions in statistics-only SurfaceEnergy calls. Replaced independent axis sorting with a conservative remap of folded energy distributions. Added stable evaluation of Eq. 3.4 for strong thermalization and a zero-abrasion-time guard. Documented the additional LISE assumption of suppressing excitation by the unthermalized fraction. Numerical and call-path regression tests passed.

18.8.13 09/15/26

AA Thermalization now is available in calculations (it was disable in 17.8.11) Correction in the final energy cell in the case of thermalization

18.8.14 09/16/26

Experimental pre-equilibrium research framework, Stage 1. Default-Off legacy isolation; persistent kernel metadata; exact Identity Kernel on raw excitation; normalized full-domain abrasion priors and LISE statistical decay probabilities. Report final-state support L_f , conditioned posteriors, and numerical omissions separately. Supplied reader pending.

18.8.16 09/17/26

Experimental pre-equilibrium research settings UI. Added a default-Off master checkbox to the prefragment excitation dialog and a dedicated Identity/Supplied kernel settings dialog. Supplied Kernel remains explicitly blocked pending its reader, with no silent fallback.

18.8.17 09/17/26

Experimental pre-equilibrium supplied-kernel integration. Added strict table validation, pinned checksums, LISE mass/reaction signatures, exact residual-energy input, and Identity template export. Research reports retain dimensionless joint contributions, final-state support L_f , posterior probabilities and numerical-loss diagnostics. Kernel metadata round-trips Unicode/multiline text. Invalid inputs stop without a fallback. Application, parser, UI and production tests passed before this version bump; all 12 legacy reference cases matched exactly. Version increment not compiled.

B2 Detailed version ledger

18.8.18 09/18/26

Experimental pre-equilibrium forward research architecture. Added single-prefragment and normalized abrasion-ensemble calculations: prefragment -> Identity or supplied kernel -> existing LISE statistical decay. Added exact residual-energy propagation, merged-state decay caching, complete dimensionless final-state support L_f , joint conditioned posteriors, parent and multiplicity marginals, and explicit omitted-probability diagnostics. The main research UI now follows the forward path; the final-isotope parent search remains available as an optional diagnostic. Supplied tables are tied to their declared forward domain and are rechecked before results are shown. Release build, forward/ensemble, supplied-kernel and UI tests passed after the final reporting fixes. All 12 Off-mode legacy cases matched exactly. Research terminal fission and statistical-breakup probabilities now use the final LISE channel values after energy-distribution rescaling; heavy-nucleus residue + fission + breakup + numerical-loss balances close to roundoff. The forward-ensemble UI reports its abrasion-parent count before execution. U-238 benchmarks at the 64-point grid took 0.14 s for 2 parents, 0.74 s for 5 parents and 3.57 s for 9 parents on the development system. Identity-versus-supplied validation now records distinct L_f and fast-neutron/ proton probabilities for a conserved synthetic table. Research settings explicitly state that On does not yet modify ordinary LISE production or cross sections.

18.8.19 09/25/26

Experimental pre-equilibrium forward validation and probability accounting. Reject unexplained statistical probability deficits as well as excesses; numerical omissions must be explicitly accounted for. Validate cached decay maps before use and skip statistical decay for zero-probability branches. Changing the forward abrasion domain or conditioned isotope invalidates the displayed research result and disables export until recalculation. Corrected the synthetic Identity/Supplied comparison to use the same input and statistical response: $L_f = 0.50 / 0.40$, respectively. The earlier 0.37 Identity comparison used a different response and was not a controlled test. Added regression tests and documented remaining convergence, physical-kernel and interactive-UI validation limits. Ordinary AA production is unchanged. Focused tests and production integration passed before this version bump; all 12 legacy reference cases matched exactly. Version increment not compiled.

B3 Detailed version ledger

18.8.20 09/27/26

Experimental pre-equilibrium numerical diagnostics and adaptive benchmarks. Added fixed-input grid studies and temporary instrumentation of statistical excitation remapping. Identified threshold-interpolation and synthetic-tail artifacts; independent conservative transport remains a diagnostic only. Added complete fast-2T point-chain references through 8192 input cells and adaptive interior-Gauss integration with explicit probability accounting. Analytic tests cover channel thresholds, invalid probabilities and limits. Low-energy Pb-206/U-236 adaptive refinement changed isotope support by $L1 = 8.59e-6 / 1.78e-6$; tighter runs took 0.90 / 0.73 s on this system. These refinement indicators are not certified convergence/error bounds. All six high-energy synthetic Re/Pb/U integrations exceeded their 5 s budgets; they report incomplete and export no partial isotope probabilities. Documented runtime limits and remaining backend work in `VALIDATION_NOTES.md`. No production decay method, pre-equilibrium physics or legacy policy changed. Analytic and integration checks passed before this version bump, including all 12 exact legacy comparisons. Version increment not compiled.

18.8.21 10/01/26

Experimental pre-equilibrium probability transport and convergence studies. Profiled fast-2T point trees; exact-state cache hits were zero in the tested cases. Added diagnostic cache-free and bounded traversals with explicit omitted path probability; retained outcomes are never renormalized. Added a separate conservative distribution prototype that accumulates probability and mean excitation per nucleus/energy bin. All 30 initial Re/Pb/U benchmark cases completed; fine broad-input runs took 5.3..6.9 s. Added optional research bin phases and partitions at existing LISE channel openings. All 48 low-energy phase/refinement cases passed probability checks; threshold splitting helps some cases but does not establish general accuracy. Separated Pb input quadrature from transport error in 21 matched-reference cases. Fine transport agrees within $9e-7$ isotope L1 for the tested inputs; the underlying input quadrature still changes appreciably with refinement. Added eight independent Pb quadrature checks: uniform references through 131072 cells and adaptive seeds 64/96/128 at tolerances $1e-6$ and $1e-7$. Adaptive seed spread remains about $5e-7$ despite the $1e-7$ indicator; neither that indicator nor probability closure is treated as a certified error bound. New backends remain isolated research prototypes, outside application dispatch. No new pre-equilibrium emission law, ABRABLA physics or production policy. Integration checks passed, including all 12 exact legacy comparisons. Detailed methods and limitations: `pre_equilibrium/CONSERVATIVE_PROTOTYPE.md`. Version metadata updated after testing; application executable not rebuilt.

B4 Detailed version ledger

18.8.22 10/02/26

Research-only Pb quadrature tracing attributes adaptive seed disagreement to exact shared energy panels and individual isotope supports. Largest discrepancy: Pb-201/Pb-200 at 57.000..57.10546875 MeV; local indicators substantially underestimate the measured difference. No physics policy or production dispatch changed. Three traced runs reproduce prior maps exactly; analytic checks and all 12 historical legacy comparisons pass.

18.8.23 10/04/26

Added 24 research-only fixed-rule Pb panel integration checks. Found false convergence when sampling misses sharp isotope-support transitions near panel edges. Saved point maps bracket Pb/Hg transitions; no production threshold or physics changed. All probability checks and 12 exact legacy comparisons pass. Explicit transition splitting is the next validation step.

18.8.24 10/04/26

Added 24 transition-split Pb research integration checks. Both fixed rules converge in four suspect panels; unresolved bracket mass is reported separately ($1.03\text{e-}9$ total). Replacing these panels reduces adaptive seed disagreement about 28-fold; remaining error is not certified. Conservation and 12 legacy regressions pass. No production calculation changed.

18.8.25 10/04/26

Added research-only automatic support-transition discovery and analytic limitation tests. Two Pb grids find 55 brackets, but miss an Hg probability transfer with unchanged isotope support. Twelve refinement cases show discovery-grid agreement is not an accuracy guarantee; probability-value jump detection remains required. Conservation and 12 legacy checks pass. No production physics or calculation dispatch changed.

18.8.26 10/04/26

Added research probability-value transition detection, catching the Hg transfer missed by support-only discovery. Twelve Pb refinement cases reduce finest step difference about 1400-fold to $2.21\text{e-}10$ isotope L1. One unmatched bracket remains under investigation; no certified global accuracy claim. Analytic, conservation and 12 legacy checks pass. No production calculation or physics parameters changed.

B5 Detailed version ledger

18.8.27 10/04/26

Resolved research detector miss of a small Au-199/Au-198 transfer: relative gating hid the jump beneath smooth variation. Corrected diagnostic mode separates refinement/acceptance floors and disables the relative gate. Eight corrected cases find 96 matching brackets; floors $1e-8/1e-9$ give identical maps. Analytic, conservation and 12 legacy checks pass. A tighter $1e-10$ run exceeded its budget and is recorded as incomplete. No production calculation or physics parameters changed.

18.8.28 10/05/26

Extended research detector validation to Re-200 and U-236. Re maps remain unchanged. Quadratic comparison resolves U discovery budget exhaustion; both grids complete with 129 matching brackets. U fission/residue integration still needs refinement (isotope L1 change $3.25e-6$). Analytic, conservation and 12 legacy checks pass; no production physics or application calculation dispatch changed.

18.8.29 10/05/26

U research refinement locates the remaining residue/fission integration difference primarily at 28.30..29.47 MeV. Eight runs complete; finest refinement changes isotope support by $6.56e-8$ (not a certified error bound). Interval accounting and 12 legacy checks pass. Next work targets that interval with reduced test scope. Production calculations remain unchanged.

18.8.30 10/05/26

Resolved the localized U-236 research quadrature problem to a steep, continuous residue/fission transfer near 28.338615 MeV. A nested dense reference converges to $1.06e-13$ isotope L1, and diagnostic replacement of that panel removes the saved step dependence to roundoff without renormalization. A general vector adaptive-Simpson prototype reproduces the dense map within $3.10e-13$ using at most 185 point decays for the tested panel. This remains research-only numerical infrastructure; finite sampling is not a completeness proof and production dispatch is unchanged. Focused probability/conservation checks passed. Per the reduced test scope, no application build or new full legacy-regression run was performed.

18.8.31 10/05/26

Combined research-only probability-transition brackets with vector adaptive-Simpson integration over the retained full U-236 excitation domain. Both independent grids complete with explicit unresolved mass and probability balance. Tightening tolerance $1e-8 \rightarrow 1e-9$ reduces cross-grid isotope L1 from $1.44e-8$ to $1.19e-9$ and fission difference to $5.14e-10$. Different unresolved bracket masses remain explicit and unnormalized. This is a measured consistency result, not certified convergence. Production dispatch and physics remain unchanged. No LISE application build or additional legacy-regression run was performed.

B6 Detailed version ledger

18.8.32 10/05/26

Moved transition discovery and adaptive probability integration from the diagnostic harness into reusable pre-equilibrium research modules. The selected-prefragment Identity action now uses this method only when the master option is On. Raw cubic-interpolated excitation is normalized once before decay; transition mass, statistical omission, integration indicator and numerical controls remain explicit without survivor renormalization. The multi-parent ensemble retains its existing research grid pending runtime optimization. Pure numerical and focused LISE On-path checks pass; Off is rejected by the adapter. The LISE application and full legacy suite were not rebuilt in this reduced test pass.

18.8.33 10/06/26

Research runtime comparisons retain the 4096-cell default: reducing seeds, omitting small upper tails, and direct traversal did not resolve Si-30 timeouts. Tail omission defaults to zero and the tested recursive evaluator remains connected. Cancellation now reaches individual statistical-decay trees. Incomplete benchmark reports no longer claim a successful numerical result. Ordinary LISE dispatch is unchanged; application not rebuilt.

18.8.34 10/06/26

Added isolated Si-30 point profiling. The bounded 30 s discovery run completed 75091 evaluations through about 82.25 MeV; individual points were inexpensive (slowest 4.55 ms). Dense discovery refinement dominates this trace, rather than a few expensive point trees. The run is incomplete and supplies no converged isotope map. Research defaults and physics are unchanged; no application build or full legacy regression in this pass.

18.8.35 10/06/26

Reduced allocation overhead in research probability-vector comparisons without changing samples or results. Si-30 trigger sweeps through $1e-4$ remained incomplete and produced no probability maps; relaxed discovery was not promoted. The user-facing single-prefragment Identity path now uses LISE's bounded normalized excitation grid with existing ManagementAA decay. The high-resolution adaptive point integrator remains a diagnostic. The Off path and statistical physics are unchanged; application not rebuilt.

18.8.36 10/07/26

Improved selected-prefragment research report formatting: excitation and $F(q)$ use compact fixed precision, probability diagnostics and isotope rows use scientific notation, and runtime values use concise significant-digit formatting. Calculations and stored probability values remain full precision.

B7 Detailed version ledger

18.8.37 10/07/26

Extended compact research formatting to the forward-ensemble/posterior and optional backward-parent reports. Exact zero/one probabilities remain 0/1; other probability columns use four-decimal scientific notation, while F(q) context uses three fixed decimals. Numerical calculations are unchanged.

18.8.38 10/07/26

Added mode-specific Evaporation Calculator sizing. Pre-equilibrium research mode enforces a 1200x700 minimum for its controls and output, while returning to Off restores the preceding regular-mode minimum and window size.

18.8.39 10/07/26

Reflowed pre-equilibrium controls from oversized horizontal rows into a compact multi-row layout. Research mode now uses an 1100x680 minimum; Off still restores the preceding regular-mode size and minimum constraint.

18.8.40 10/07/26

Added a clearly labeled synthetic one-neutron validation-kernel generator for the selected prefragment. It creates and pins a complete matching-grid table with a 20% algebraic 1n branch above threshold, exact LISE-mass energy balance, and two-body kinetic/recoil partition. It is explicitly not a physical emission model and does not affect Legacy Reference calculations. Changed objects and the production integration suite pass, including the selected-forward x=1,y=0 path and all 12 exact legacy comparisons. The application executable was not linked for this version increment.

18.8.41 10/07/26

Fixed the synthetic supplied-kernel forward validation path exceeding the recursive point-decay state limit. Its high-energy Identity contribution is now accumulated on the normalized LISE excitation grid and passed once to ManagementAA; the one-neutron Er=0 contribution remains exact. Kernel-node probabilities and the conservation ledger are unchanged. The affected MSVC object and production integration suite pass, including the generated x=1, y=0 channel and all 12 exact legacy comparisons. Ordinary AA is unchanged.

B8 Detailed version ledger

18.8.42 10/07/26

Applied compact user-facing formatting to selected-prefragment Supplied results. Total residue probability uses five fixed decimals; fission, statistical breakup and numerical-balance diagnostics use four-decimal scientific notation. Channel and isotope probability tables use four-decimal scientific notation while retaining exact 0 and 1. Calculations are unchanged.

18.8.43 10/07/26

Extended the generated synthetic one-neutron validation kernel to the full normalized forward abrasion ensemble. A reusable bounded evaluator now serves both selected-prefragment and ensemble paths. Added a UI action that creates, pins and selects a complete synthetic table for the displayed abrasion domain. Ensemble validation checks both multiplicity channels, dimensionless L_f and the normalized conditioned posterior. All 12 exact legacy comparisons pass. The algebraic 20% branch remains software validation only; ordinary AA and general supplied-kernel physics are unchanged.

18.8.44 10/07/26

Applied compact display formatting to synthetic-kernel metadata and forward runtime. $P(1n)$ is shown with three fixed decimals, including for existing saved validation tables, and runtime uses four significant digits. Newly generated tables store the compact provenance. Calculations are unchanged.

18.8.45 10/07/26

Added a second algebraic validation fixture with separate $1n$, $1p$ and correlated $1n1p$ branches. Default test masses are 0.100, 0.100 and 0.050 when energetically allowed; unavailable mass remains explicitly in Identity. Exact LISE masses enforce A , Z and energy conservation, with $E_r=0$ emitted branches and explicit total kinetic/recoil energy. Selected-prefragment and full-ensemble UI generators are provided. Production integration verifies all four multiplicity channels, L_f /posterior accounting and all 12 exact legacy comparisons. This remains software validation, not a physical model.

18.8.46 10/07/26

Added a user-facing diagnostic when supplied fast-emission branches exist but none contributes to the selected final isotope, explaining the resulting Identity-only posterior and the generated $E_r=0$ validation-domain remedy. Production tests now condition separately on branch-accessible residuals and require nonzero $1n$, $1p$ and $1n1p$ posterior contributions. The affected MSVC object, production integration suite and all 12 exact legacy comparisons pass.

B9 Detailed version ledger

18.8.47 10/07/26

Added a finite-residual-excitation algebraic validation fixture. Allowed 1n, 1p and 1n1p branches retain $E_r=40$ MeV; exact LISE masses define Delta-M and all remaining energy is assigned to fast kinetic energy plus recoil. Selected and forward-ensemble UI generators feed these residuals into the existing LISE statistical decay. Tests prove downstream de-excitation, including a deterministic lighter-residue outcome, probability closure and all 12 exact legacy comparisons. The fixture is not a physical emission model.

18.8.48 10/07/26

Added reusable supplied-kernel diagnostics before statistical decay. Selected and ensemble reports now distinguish unconditional $U(x,y)$ from final-isotope posteriors and show input/residual energy ranges, mean $E_r/K_{fast}/K_{recoil}$ by multiplicity, kernel probability closure, node normalization error and the maximum complete energy-ledger residual. Production tests verify unit kernel probability and sub-tolerance conservation; all 12 legacy comparisons pass.

18.8.49 10/07/26

Added a pre-run supplied-kernel inspector in the research panel. It identifies whether the pinned table matches the selected-prefragment or displayed forward ensemble domain, performs complete signature/mass/probability/energy validation and displays unconditional diagnostics without running statistical decay. Tables matching neither domain fail explicitly. Selected and ensemble reports retain diagnostic coverage; production and all 12 legacy comparisons pass.

18.8.50 10/07/26

Added reaction-matched editable user-kernel templates and a supplied-format guide defining metadata, branch fields, Delta-M energy conservation and all reported probability quantities. Kernel errors now identify the parent, excitation node, branch and energy residual. Forward reports distinguish input nodes, branch records and merged statistical states. Focused and production tests pass, including all 12 exact legacy comparisons.

18.8.51 10/07/26

Added a built-in supplied-kernel format guide to the research panel so the exact export/edit/select/inspect workflow, branch fields, Delta-M ledger and probability definitions remain available in installed builds. The guide states that results are dimensionless and that no ABRABLA physics is used.

B10 Detailed version ledger

18.8.52 10/08/26

Added forward-ensemble performance diagnostics for scaling studies. Identity and Supplied reports now show preparation, statistical-decay and total runtime plus aggregate input states, merged states, decay evaluations, cache hits and calculated nuclei. Probability calculations and legacy dispatch are unchanged.

18.8.53 10/09/26

Shared identical emitted-residual statistical-decay states across abrasion parents in generated validation ensembles. Each parent's continuous Identity distribution remains isolated from point-state caching. Reports now label supplied input-state totals as branch records. Probabilities are unchanged.

18.8.54 10/09/26

Added monotonic phase-scaled progress across supplied-table preparation and every abrasion parent's kernel/statistical-decay work. Large research ensembles now remain responsive and cancellable within a parent calculation; cancellation still returns no completed result. Physics is unchanged.

18.8.55 10/10/26

Added a pre-run emitted-state workload forecast to supplied-kernel inspection: positive and emitted branch records, parent-merged emitted states, globally unique emitted decays and potential cross-parent cache reuse. The forecast excludes Identity distributions and runs no statistical decay.

18.8.56 10/10/26

Added post-run scaling metrics to forward-ensemble reports: cache-reuse percentage and statistical-phase wall time per executed decay. These values are diagnostics only and do not impose pruning, cutoffs or runtime policy.

B11 Detailed version ledger

18.8.57 10/10/26

Added the reproducible Ca/Pb/U Identity and finite-E* validation-kernel scaling matrix, froze supplied-kernel interchange contract version 1, and recorded the non-ABRABLA physics requirements for a future physical kernel.

18.8.58 10/10/26

Added cancellable selected-parent and forward-ensemble provider-package exports with normalized abrasion/input probabilities and the complete reachable LISE nuclear-mass inventory. No emission law or cutoff is added.

18.8.59 10/10/26

Added a strict provider-exchange round-trip validator. It reconstructs the exported normalized context and complete mass domain, rejects altered or incomplete requests, and validates returned kernels before statistical decay without repair, renormalization or an internal emission model.

18.8.60 10/10/26

Added the standalone LISEKernelValidator console target for external kernel providers. It checks exported request/response pairs with the production contract and reports checksums and structural counts without statistical decay, cross sections or changes to ordinary LISE calculations.

18.8.61 10/10/26

Added research-panel validation of an exported provider request together with its returned kernel. The action reports checksums and exact structural coverage without selecting the kernel, changing settings, invalidating a completed result, running statistical decay or calculating cross sections.

B12 Detailed version ledger

18.8.62 10/10/26

Moved the experimental workflow into a dedicated modeless pre-equilibrium research window with grouped Calculations, Kernel exchange and Software validation tabs. The Evaporation Calculator retains its normal layout, size, modality and legacy controls while research mode is enabled.

18.8.63 10/10/26

Moved the compact pre-equilibrium master, kernel, settings and research launch strip to the bottom of the Evaporation Calculator, directly above the existing file-output and Help/Quit controls. Research behavior is unchanged.

18.8.64 10/10/26

Temporarily froze the experimental pre-equilibrium user interface. Its controls remain visible but disabled in the Evaporation Calculator and the Excitation Energy dialog. Loaded On states are forced to Legacy Reference while inactive kernel metadata and all research implementation are retained for later reactivation.

C Evidence register and references

Earlier project documents

1. doc/LISE_pre_equilibrium_schemes.docx. Original complete/kernel schemes, conservation ledger and support/posterior definitions. Retained as historical source.
2. doc/LISE_PreEquilibrium_Architecture_and_Status_18_8_21.docx. Forward architecture, completed infrastructure, prototype numerical status and proposed next steps as of 1 October 2026.

Source and development records

- L_Init/lise_version.h: source snapshot 18.8.91, last revision 17 October 2026; pre-equilibrium development entries 18.8.12–64.
- o_Evap/pre_equilibrium/README.txt and VALIDATION_NOTES.md: implementation contract, historical experiments, limitations and recorded validation outcomes.
- o_Evap/pre_equilibrium/KERNEL_CONTRACT_V1.txt and SUPPLIED_KERNEL_FORMAT.txt: frozen version-1 exchange schema, normalization and conservation rules.
- o_Evap/pre_equilibrium/KERNEL_SCALING.txt: measured 18-case Ca/Pb/U timings and provider-package inventories.
- o_Evap/pre_equilibrium/PHYSICAL_KERNEL_DECISION.txt: recorded decision to require a defensible dynamical initial state before adopting an internal physical kernel.
- o_Evap/pre_equilibrium/CONSERVATIVE_PROTOTYPE.md: conservative transport methods and numerical evidence retained from the prototype phase.
- o_Evap/o_PreEquilibriumSettings.h and both public dialog implementations: verified current freeze and preserved metadata behavior.

Physics literature named in the project record

J. J. Gaimard and K. H. Schmidt, Nuclear Physics A531 (1991), 709. Eq. 3.4 is the source associated with the existing LISE thermalization factor. The LISE implementation comment distinguishes equilibration from its additional phenomenological suppression assumption.

J. J. Griffin, Statistical Model of Intermediate Structure, Physical Review Letters 17, 478 (1966). DOI 10.1103/PhysRevLett.17.478.

C. Kalbach, Two-component exciton model Basic formalism away from shell closures, Physical Review C 33, 818 (1986). DOI 10.1103/PhysRevC.33.818.

R. Nepeivoda et al., Pre-Equilibrium Clustering in Production of Spectator Fragments in Collisions of Relativistic Nuclei, Particles 5, 40–51 (2022). DOI 10.3390/particles5010004.

The latter references identify model families discussed in the existing decision record. No parameters from them are instantiated in the present kernel, and no new literature review or original-paper verification is claimed by this documentation update.